

(12) UK Patent Application (19) GB (11) 2 393 540 (13) A

(43) Date of A Publication 31.03.2004

(21) Application No: 0317525.4
 (22) Date of Filing: 25.07.2003
 (30) Priority Data:
 (31) 10216437 (32) 12.08.2002 (33) US

(71) Applicant(s):
 Hewlett-Packard Development Company L.P. a limited
 partnership organised under the laws of Texas
 U.S.A., 20555 S.H.249, Houston,
 Texas 77070, United States of America

(72) Inventor(s):
 Andrew H. Barr
 Ricardo Espinoza-Ibarra

(74) Agent and/or Address for Service:
 Carmichael & Ransford
 43 Bloomsbury Square, LONDON,
 WC1A 2RA, United Kingdom

(51) INT CL⁷:
 G06F 1/32

(52) UK CL (Edition W):
 G4A APWR

(56) Documents Cited:
 US 20020007463 A1

(58) Field of Search:
 UK CL (Edition W) G4A
 INT CL⁷ G06F, H05K
 Other: WPI, EPDOC, PAJ, INSPEC, IBM TDB, IEEE
 Xplore, Internet

(54) Abstract Title: **Managing the operating frequency of blades**

(57) In a system, method and apparatus for managing the operating frequency of blades in a blade-based computer system (100) based upon performance requirements, a first blade (120) that requires a specific power allocation is operated at a specific frequency and consumes a portion of the thermal and power budget for the blade-based computer system. A second blade (130) that requires a specific power allocation is operated at a specific frequency and consumes a portion of the thermal and power budget for the blade-based computer system (100). Additionally, the overall power and thermal budget of the blade-based computer system (100) is maintained.

500

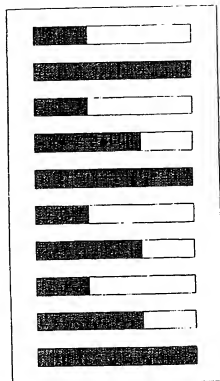


Figure 9

GB 2 393 540 A

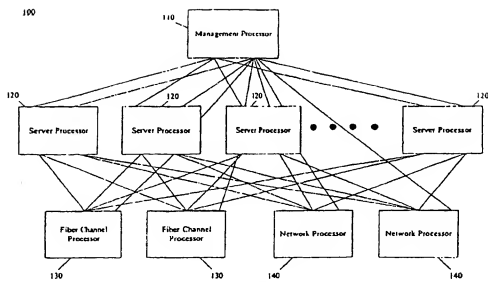


Figure 1

200

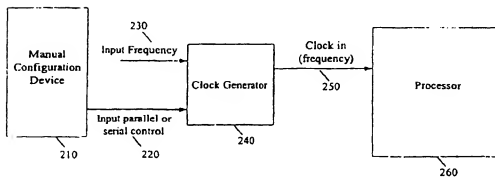


Figure 2

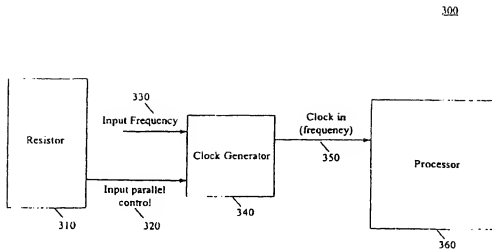


Figure 3

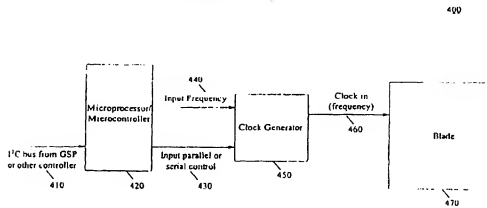


Figure 4

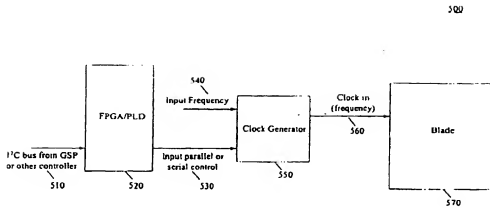


Figure 5

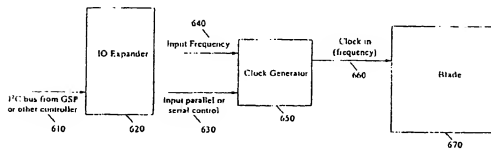


Figure 6

700

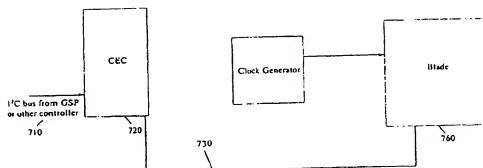


Figure 7

800

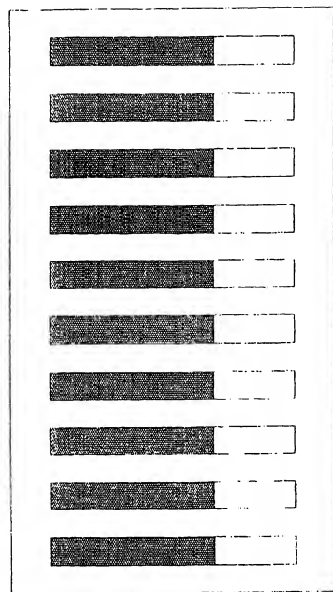


Figure 8

9002

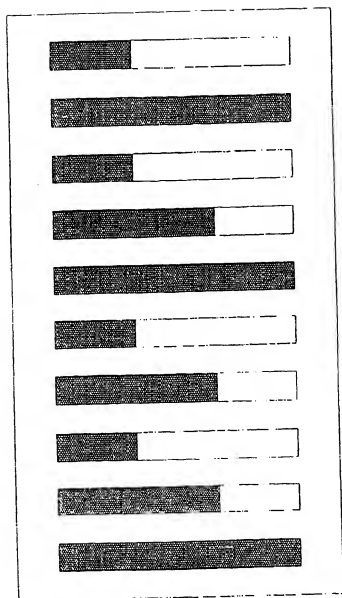


Figure 9

**SYSTEM, METHOD AND APPARATUS FOR THE FREQUENCY
MANAGEMENT OF BLADES IN A BLADED ARCHITECTURE BASED ON
PERFORMANCE REQUIREMENTS**

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202607-1), entitled "SYSTEM AND METHOD FOR THE FREQUENCY MANAGEMENT OF COMPUTER SYSTEMS TO ALLOW CAPACITY ON DEMAND" to Andrew H. BARR, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202610-1), entitled "SYSTEM, METHOD AND APPARATUS FOR PERFORMANCE OPTIMIZATION AT THE PROCESSOR LEVEL" to Ricardo ESPINOZA-IBARRA, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202612-1), entitled "SYSTEM AND METHOD FOR LOAD DEPENDENT FREQUENCY AND PERFORMANCE MODULATION IN BLADED SYSTEMS" to Ricardo ESPINOZA-IBARRA, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202878-1), entitled "VOLTAGE MANAGEMENT OF BLADES IN A BLADED ARCHITECTURE BASED ON PERFORMANCE REQUIREMENTS" to Andrew H. BARR, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202880-1), entitled "VOLTAGE MODULATION IN CONJUNCTION WITH PERFORMANCE OPTIMIZATION AT PROCESSOR LEVEL" to Andrew H. BARR, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202881-1), entitled "SYSTEM AND METHOD FOR MANAGING THE OPERATING FREQUENCY OF PROCESSORS OR BLADES" to Ricardo ESPINOZA-IBARRA, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202882-1), entitled "SYSTEM AND METHOD FOR MANAGING THE OPERATING FREQUENCY OF BLADES IN A BLADED-SYSTEM" to Ricardo ESPINOZA-IBARRA, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100202916-1), entitled "VOLTAGE MANAGEMENT OF PROCESSORS IN A BLADED SYSTEM BASED ON LOADING" to Andrew H. BARR, et al.; U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100203096-1), entitled "SYSTEM AND METHOD FOR VOLTAGE MANAGEMENT OF A PROCESSOR TO OPTIMIZE PERFORMANCE AND POWER DISSIPATION" to Andrew H. BARR, et al., and U.S. Patent Application Serial No. 10/---,--- (Attorney Docket No. 100203638-1), entitled "MANAGEMENT OF A MEMORY SUBSYSTEM" to Andrew H. BARR, et al., all of which are concurrently

being filed herewith under separate covers, the subject matters of which are herein incorporated by reference.

BACKGROUND

Bladed servers, or bladed computer systems, are computing systems that provision servers or other computer resources on individual cards, or blades. There are many types of blades - processor blades, server blades, storage blades, network blades, etc. - and one or more of each type of blade is typically housed together in a single structure, thus creating high-density computing systems with modular architectures, ensuring flexibility and scalability of the systems, and reducing space requirements. Server blades, along with storage, networking, and other blades, are typically installed in a rack-mountable enclosure, or chassis, which hosts multiple blades. The multiple blades share common resources such as cabling, power supplies, and cooling fans.

The telecommunications industry has been using blade server technology for many years. More generally, the condensed blade server architecture benefits people and businesses that: (1) use the Internet to generate revenue and to provide services to customers; (2) are moving some of their business processes to the Web; and/or (3) need the flexibility to deploy Internet-edge applications in their own data center. Because of recent developments in technology, blade servers are now used for applications such as Web hosting, Web caching, and content streaming.

In Web caching applications, frequently-requested Web content is stored closer to the user, thus allowing for quicker retrieval of objects by the user and reducing the time and bandwidth required to access the Internet. Since companies and individuals are now streaming media (e.g., video and audio) to more effectively communicate both internally and externally, a massive growth of rich media content delivery on the Internet has occurred. Bladed servers are being used to meet the new demands created as a result of this growth.

Though bladed servers provide many advantages, several engineering challenges arise when using bladed servers. Among these challenges is the challenge of designing and operating a bladed system such that sufficient heat is dissipated in the limited space available in the chassis that hosts the system. To address heat dissipation challenges, bladed server systems are designed within an underlying power and thermal envelope. For example, when a chassis that hosts a bladed system has a limited amount of airflow available to cool the blades (i.e., when the system can only dissipate a limited amount of heat), then the chassis is designed for a limited amount of power consumption and an

1 associated limited performance of the blades. Some known power limiting strategies
2 include powering down a CPU functional unit, e.g., a floating-point unit or an on-die
3 cache, or trading off speed for reduced power consumption in a hard drive.

4 In a system where the ability of the chassis to cool is limited to X and there are Y
5 blades, each blade can only approximately contribute X/Y to the dissipated power in the
6 chassis. Thus, each blade is limited to the performance associated with an X/Y power
7 level.

8 In a bladed architecture, multiple blades, each representing a separate system, are
9 present in the same chassis. Associated with the chassis are a specific set of power and
10 thermal requirements. Specifically, these requirements put a limit on the amount of power
11 that can be consumed by the blades. This power limitation puts a limitation on the
12 frequency that the processors on the blade can run, and thus limits the performance.

13 Prior solutions included running all the blades at a performance level less than the
14 blades' maximum in order to meet the overall chassis power and thermal cooling budget.
15 Thus, the performance of each blade is limited to fall within these budgets.

16 SUMMARY

17 In one embodiment of the frequency management of blades in a bladed
18 architecture based on performance requirements, a first blade that requires a specific
19 power allocation is operated at a specific frequency and consumes a portion of the
20 thermal and power budget for the blade-based computer system. A second blade that
21 requires a specific power allocation is operated at a specific frequency and consumes a
22 portion of the thermal and power budget for the blade-based computer system.
23 Additionally, the overall power and thermal budget of the blade-based computer system is
24 maintained.

25 In another embodiment of the frequency management of blades in a bladed
26 architecture based on performance requirements, a first blade that requires a specific
27 power allocation is operated at a specific frequency and consumes a portion of the
28 thermal and power budget for the blade-based computer system. A second blade that
29 requires a specific power allocation is operated at a specific frequency and consumes a
30 portion of the thermal and power budget for the blade-based computer system.
31 Additionally, a means to maintain the frequency level of at least the first blade or second
32 blade is utilized.

DESCRIPTION OF THE DRAWINGS

The detailed description will refer to the following drawings, wherein like numerals refer to like elements, and wherein:

FIGURE 1 shows a block diagram depicting one embodiment of the basic modular building blocks of a bladed architecture system;

FIGURE 2 illustrates a block diagram depicting one methodology of managing the operating frequency of individual blades by use of a manual configuration device;

FIGURE 3 illustrates a block diagram depicting another method of managing the operating frequency of individual blades by use of resistors;

FIGURE 4 illustrates a block diagram depicting another method of managing the operating frequency of individual blades by use of a microcontroller or microprocessor;

FIGURE 5 illustrates a block diagram depicting another method of managing the operating frequency of individual blades by use of a field-programmable gate array (FPGA) or programmable logic device (PLD);

FIGURE 6 illustrates a block diagram depicting another method of managing the operating frequency of individual blades by use of an I/O-Expander chip (I/OX);

FIGURE 7 illustrate therein a block diagram depicting the methods of managing the operating frequency of individual blades in an IA based architecture system;

FIGURE 8 illustrates a block diagram depicting a series of blades inside of a bladed architecture chassis operating at the same frequency level; and

FIGURE 9 illustrates a block diagram depicting a series of blades inside of a bladed architecture chassis operating at the different frequency levels.

DETAILED DESCRIPTION

The preferred embodiments of the frequency management of blades in a bladed architecture based on performance requirements will now be described in detail with reference to the following figures, in which like numerals refer to like elements. With reference to FIGURE 1 of the Drawings, there is illustrated therein a block diagram depicting one embodiment of the basic modular building blocks of a bladed architecture system, as generally designated by the reference numeral 100. A management blade 110 supervises the functions of the chassis and provides a single interface to the consoles of all the servers installed. As shown in FIGURE 1, server blades 120 are in communication with the management blade 110. The server blades 120 are, in turn, in communication with other blades that perform specific functions. For example, as seen in FIGURE 1, server blades 120 are in communication with fiber channel blades 130 and network blades

1 140. It is to be appreciated that the various blades in a bladed architecture system may be
2 processor blades, server blades, network blades, storage blades or storage interconnect
3 blades, etc.

4 As discussed above, it is desirable to create a bladed architecture system in which
5 each blade may run at a different frequency even though the blades may have the same
6 exact hardware and software. Blades that host applications that require higher
7 performance are allowed to run at an increased frequency, and thus consume more of the
8 chassis thermal and power budget. Blades that host applications that require a lower level
9 of performance are run at a lower frequency, and thus consume less of the chassis thermal
10 and power budget. As a result, the overall thermal and power requirements are still met
11 on average, yet a more optimal performance solution is reached. Further, as the relative
12 performance requirements on the blades change, the frequencies for the blades can be
13 changed, thus changing the performance and thermal/power budget allocation. It is to be
14 appreciated that the apparatus and methods disclosed herein can be applied to PA, IA-32,
15 IA-64 or any processor blade technology.

16 FIGURES 2-7 illustrate various methods for managing the operating frequency of
17 the various processors or blades based upon performance requirements. In some bladed
18 architectures, the processor core frequency of the CPU is asynchronous to the bus
19 operating frequency; therefore, a change to the processor core frequency can be made
20 independently of the system bus frequency. Currently, Hewlett Packard uses the SPHYR-
21 T ASIC[®] as a frequency-synthesizer for generating the processor clock of the PA-RISC[®]
22 systems. However, one skilled in the art would readily recognize that other synthesizers
23 may also be used. In other embodiments of processor architectures, e.g., IA-32 and IA-
24 64, the processor core frequency of the CPU is not necessarily asynchronous to the bus
25 operating frequency. In IA processors the processor core frequency runs at a
26 programmable multiple of the bus operating frequency. Managing the operating
27 frequency of individual blades can be applied to both architectures by modifying the
28 output of the clock chip, and thus, setting the processor frequency accordingly.

29 Generally, frequency synthesizer chips used to generate the processor's clocks
30 have parallel or serialized configuration bits that allow one to choose the ratio of the input
31 clock to the output clock (synthesized frequency). This allows one to run the processors
32 at a different frequency upon reboot of the blade. The frequency synthesizer typically has
33 an input frequency from a core crystal. Through controlling the serial or parallel pins, the
34 frequency synthesizer provides the output frequency ratio that is sent to the processors.

One skilled in the art would also appreciate that if multiple frequency synthesizers are utilized, i.e., two (for a 2-way system) or N (for an N-way system), to generate the processor's clocks, the processors are able to operate at different frequencies within the range of frequencies supported by the processor. FIGURES 2-6 illustrate different methodologies of the how the ratio pins of the different frequency synthesizers can be controlled. It is to be appreciated the below described methodologies are used to modulate the frequency at the blade system level.

With reference now to FIGURE 2 of the Drawings, there is illustrated therein a block diagram depicting one methodology of managing the operating frequency of individual blades by use of a manual configuration device 210, as generally designated by the reference numeral 200. As seen in the figure, the manual configuration device 210 is added to a readily accessible part of the system. One of ordinary skill in the art would recognize that there are many common manual configuration devices that are capable of performing the desired function, e.g., dip switches, jumpers installed over pin headers, rotational configuration switches, and solder bridges, etc. Thus, the operator is allowed to set the frequency of the processors upon reboot of the unit, based on predetermined performance requirements. The input frequency 230 and signal 220 from the manual configuration device 210 are used to generate an output frequency 250 at the clock generator, or frequency synthesizer, 240 that is used by the processors for the blade 260. It is to be appreciated that this method of changing the frequency of the blade by use of a manual configuration device, and the other methods described herein below, may be applied to various types of processor architectures, e.g., PA-RISC, DEC Alpha, MIPS, PowerPC, SPARC, IA-32 and IA-64.

With reference now to FIGURE 3 of the Drawings, there is illustrated therein a block diagram depicting another method of managing the operating frequency of individual blades, as generally designated by the reference numeral 300. Method 300 manages the operating frequency of individual blades by use of a resistor 310. As seen in FIGURE 3, the resistor 310 is added to a readily accessible part of the system. Thus, the operator is allowed to set the frequency of the blades upon reboot of the unit, based on predetermined performance requirements. An input frequency 330 and signal 320 from the resistor 310 are used to generate an output frequency 350 at a clock generator 340 that is used by processors on the blade 360.

With reference now to FIGURE 4 of the Drawings, there is illustrated therein a block diagram depicting the method of managing the operating frequency of individual

blades by use of a microcontroller or microprocessor 420, as generally designated by the reference numeral 400. The microcontroller or microprocessor 420 is used to interface with the user of the system to ask for the specific frequency at which each blade should run. As seen in FIGURE 4, microcontroller or microprocessor 420 receives a signal on an I²C (Inter-IC) bus 410 from a GSP (service processor) or other controller. As is known in the art, an I²C bus is a bi-directional two-wire serial bus that provides a communication link between integrated circuits. Further, a person of ordinary skill in the art would readily recognize that other control buses could perform the same functions and be substituted for the I²C bus described herein. The microcontroller or microprocessor 420 outputs a parallel or serial control 430, based upon the specific frequency designated by the user. A clock generator 450 uses an input frequency 440 and parallel or serial control 430 to generate an output frequency 460 used by the designated blade 470. The use of the microcontroller or microprocessor 420 allows the user to control the frequency synthesizers in a more transparent way than the above described register and manual configuration device methods, i.e., the user does not necessarily need to know how the settings of the configuration bits will affect the output.

With reference now to FIGURE 5 of the Drawings, there is illustrated therein a block diagram depicting the method of managing the operating frequency of individual blades by use of a FPGA (field-programmable gate array) or PLD (programmable logic device) 520, as generally designated by the reference numeral 500. As known in the art, a FPGA is a chip that can be programmed in the field after manufacture. The FPGA or PLD 520 is used to receive commands from a higher-level device, e.g., the GSP via an I²C bus 510, to control the configuration bits for the frequency synthesizer (or each synthesizer). As seen in FIGURE 5, a clock generator 550, uses an input frequency 540 and the control from the FPGA/PLD 530 to generate an output frequency 560 used by a processor or blade 570. Like the use of the microcontroller/microprocessor 420 described in FIGURE 4, the FPGA/PLD 520 allows the user to control the frequency synthesizers in a more transparent way, i.e., the user does not necessarily need to know how the settings of the configuration bits will affect the output.

With reference now to FIGURE 6 of the Drawings, there is illustrated therein a block diagram depicting the method of managing the operating frequency of individual blades by use of an I/O-Expander chip (I/OX) 620, as generally designated by the reference numeral 600. As known in the art, an I²C based I/OX 620 is an inexpensive and simple solution that can be used to transparently control each frequency synthesizer.

1 I/OX chips 620 have I/O ports, which can be forced to a particular state by writing to the
2 I/OX through an I²C command. Since I/OX typically have multiple I/O ports, it is to be
3 appreciated that one I/OX can be used to control multiple frequency synthesizers
4 individually.

5 Since I/OX chips are I²C-based, they can be controlled by any device that supports
6 an I²C interface. In Hewlett Packard's PA-RISC Blades, a logical such device would be
7 the service processor, or GSP, of the PA Blade. The user-friendly interface of the GSP
8 can be used to transparently allow the customer to control the frequency of the processors
9 without need of any low-level information, e.g., bit-settings. The input from the GSP is
10 designated in FIGURE 6 by reference numeral 610. The user at the GSP inputs
11 commands to increase the frequency of the processor. The GSP then delivers the required
12 bits to change the ratio of the frequency generator, so that the processor will run at the
13 desired level. The bit stream 610 is received by the I/OX 620. A clock generator 650
14 then uses an input frequency 640 and a control signal 630 to generate an output frequency
15 660 used by a processor or blade 670.

16 With reference now to FIGURE 7 of the Drawings, there is illustrated therein a
17 block diagram depicting the methods of managing the operating frequency of individual
18 blades in an IA based architecture system, as generally designated by the reference
19 numeral 700. As shown above, PA-based architectures have two frequency inputs that
20 are taken into the processor: the processor core input and the bus input. In IA-based
21 architectures the serial or parallel control is not modulated. IA-based architectures
22 program the processor core frequency to generate a multiple of the bus frequency. The
23 programming of the processor core frequency is done during early bus initialization using
24 multipurpose bus lines 730. This is typically controlled by the main core chipset that sits
25 on the processor bus (CEC) 720. Generally, a user interface, e.g., a GSP 710, is used to
26 program the CEC 720 to generate the appropriate control signal during early system and
27 bus initialization. The bus frequency is multiplied for a larger ratio to generate a higher
28 frequency, performance and power, or for a smaller ratio to generate a lower frequency,
29 performance and power, depending on the application at the particular blade or processor
30 760.

31 With reference now to FIGURE 8 of the Drawings, there is illustrated therein a
32 block diagram depicting a series of blades inside of a bladed architecture chassis running
33 at the same frequency, as generally designated by the reference numeral 800. The
34 shading of the individual blades indicates that each individual blade is operating at the

1 same frequency level. Operating blades at the same frequency is typical in current bladed
2 architecture systems. In addition, the shading illustrates that each blade is operating at a
3 level below the maximum level in order to remain under the maximum power allocated to
4 the system as a whole. As discussed, bladed server systems are limited by an underlying
5 power and thermal envelope. This is due to the heat produced within the blades and to
6 the limited dimensions in the chassis. When the chassis consumes a given amount of the
7 power, the chassis is typically limited in the amount of airflow that is available to cool the
8 blades. As a result, the power limitation limits the frequency that the processors on the
9 blade can run, and thus, limits the performance. The processors within the blades are thus
10 limited in their ability to operate at optimal performance and capacity because the
11 processors are configured to operate at the same frequency a frequency below their
12 maximum level.

13 With reference now to FIGURE 9 of the Drawings, there is illustrated therein a
14 block diagram depicting a series of blades inside of a bladed architecture chassis running
15 at different frequencies based upon the individual blades performance requirements, as
16 generally designated by the reference numeral 900. The shading of the figure indicates
17 the level of the frequency at which the blade is being run. As seen in the figure, blades
18 that host applications that require higher performance are allowed to run at an increased
19 frequency, and thus consume more of the chassis thermal and power budget. Blades
20 running at an increased frequency are indicated by the increased shaded area. Blades that
21 host applications that require a lower level of performance are run at a lower frequency,
22 and thus consume less of the chassis thermal and power budget. Blades running at a
23 lower frequency are indicated in the figure by the decreased shaded area. It is to be
24 appreciated that each blade may run at a different frequency even though the blades have
25 the exact same hardware and software.

26 As is known in the art, the different requirements for the different frequencies is
27 dependent upon the actual applications that each blade is performing. FIGURE 9
28 illustrates that blades with a higher power budget, and thus higher performance are run at
29 a higher frequency, and blades that have a lower power budget, and thus lower
30 performance, are run at a lower frequency, while the overall power and thermal budget
31 for the chassis is maintained. Thus, the overall thermal and power requirements are still
32 met on average, yet a more optimal performance solution is reached. Additionally, when
33 the relative performance requirements on the blades change, the frequencies can be
34 changed, thus changing the performance and thermal/power budget allocation. It is to be

1 appreciated that the principles can be applied to PA, IA32, IA64 or any processor blade
2 technology.

3 It is to be appreciated that the principles disclosed herein may be applied to a
4 system comprised of processors or blades that share a common chassis or to an
5 architecture system that spans multiple chassis. That is, the principles may be applied to
6 systems that are divided by either a physical or logical partition. For example, physically,
7 a system may include three chassis, with each chassis having eight processors. Logically,
8 the same system could be partitioned into five different web servers for five different
9 customers. Power constraints within a chassis typically concern the physical partition of
10 the system. Power constraints imposed on a customer or application that is located in
11 multiple chassis, typically concern logical partitions. One of ordinary skill in the art
12 would readily recognize that the innovations described above may be applied to both
13 physically and logically partitioned architectures.

14 While the frequency management of blades in a bladed architecture based on
15 performance requirements has been described in connection with exemplary
16 embodiments, those skilled in the art will understand that many modifications in light of
17 these teachings are possible, and this application is intended to cover any variation thereof.

18 For example, the disclosed system and method makes use of specific I²C devices
19 that are used to receive signals from an I²C bus. Other I²C devices could likewise be
20 used. Thus, the I²C devices shown and referenced generally throughout this disclosure,
21 and unless specifically noted, are intended to represent any and all devices/technologies
22 appropriate to perform the desired function. Likewise, there are disclosed several
23 processors and blades that perform various operations. The specific processor or blade is
24 not important to the disclosure. Thus, it is not applicant's intention to limit this disclosure
25 to any particular form of processor, blade or specific blade architecture.

26 Further examples exist throughout the disclosure, and it is not applicant's intention
27 to exclude from the scope of this disclosure the use of structures, materials, or acts that
28 are not expressly identified in the specification, but nonetheless are capable of performing
29 a claimed function.

1 In the claims:

2 1. A method for managing the operating frequency of blades in a blade-based
3 computer system (100), said method comprising the steps of:

4 operating a first blade (120) at a first frequency, said first blade (120) hosting
5 applications that require a first power allocation, wherein said first blade (120) consumes
6 a portion of the thermal and power budget in said blade-based computer system (100)
7 based upon said first power allocation;

8 operating a second blade (130) at a second frequency, said second blade (130)
9 hosting applications that require a second power allocation, wherein said second blade
10 (130) consumes a portion of the thermal and power budget in said blade-based computer
11 system (100) based upon said second power allocation; and

12 maintaining the thermal and power budget in said blade-based computer system
13 (100) and optimizing performance in said blade-based computer system (100) by
14 modulating at least one of said first frequency and said second frequency in said blade-
15 based computer system (100).

16 2. The method according to claim 1, wherein said first blade (120) and said second
17 blade (130) contain the same hardware and software.

18 3. The method according to claim 1, wherein said blade-based computer system
19 (100) is based on an architecture chosen from the group consisting of: PA-RISC, DEC
20 Alpha, MIPS, PowerPC, SPARC, IA-32 and IA-64.

21 4. The method according to claim 1, wherein the thermal and power budget
22 allocation is altered according to a change in the performance requirements for individual
23 blades in said blade-based computer system (100).

24 5. The method according to claim 1, wherein each blade in said blade-based
25 computer system (100) requires a unique power allocation.

26 6. The method according to claim 5, wherein each blade in said blade-based
27 computer system (100) is run at a frequency dependent upon said unique power
28 allocation.

29 7. A blade-based computer system (100) comprising:

30 a first blade (120) hosting applications that require a first power allocation, said
31 first blade (120) operated at a first frequency, wherein said first blade (120) consumes a
32 portion of the thermal and power budget in said blade-based computer system (100) based
33 upon said first power allocation;

1 a second blade (130) hosting applications that require a second power allocation,
2 said second blade (130) operated at a second frequency, wherein said second blade (130)
3 consumes a portion of the thermal and power budget in said blade-based computer system
4 (100) based upon said second power allocation; and
5 wherein the overall thermal and power budget in said blade-based computer
6 system (100) is maintained and performance in said blade-based computer system (100) is
7 optimized by modulating at least one of said first frequency and said second frequency.

8 8. The blade-based computer system according to claim 7, wherein said first blade
9 (120) and said second blade (130) contain the same hardware and software.

10 9. A system for managing the operating frequency of blades in a blade-based
11 computer (100), said system comprising:
12 a first blade (120) hosting applications that require a first power allocation, said
13 first blade (120) operated at a first frequency, wherein said first blade (120) consumes a
14 portion of the thermal and power budget in said blade-based computer (100) based upon
15 said first power allocation;
16 a second blade (130) hosting applications that require a second power allocation,
17 said second blade (130) operated at a second frequency, wherein said second blade (130)
18 consumes a portion of the thermal and power budget in said blade-based computer (100)
19 based upon said second power allocation; and
20 a management means for altering the frequency level of at least said first blade
21 (120) or said second blade (130).

22 10. The system according to claim 9, wherein said blade-based computer (100) is
23 based on an architecture chosen from the group consisting of: PA-RISC, DEC Alpha,
24 MIPS, PowerPC, SPARC, IA-32 and IA-64.



Application No: GB 0317525.4
Claims searched: 1 to 10

Examiner: Michael Powell Waters
Date of search: 19 January 2004

Patents Act 1977 : Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A		US 2002/0007463 A1 (FUNG)

Categories:

X Document indicating lack of novelty or inventive step	A Document indicating technological background and/or state of the art.
Y Document indicating lack of inventive step if combined with one or more other documents of same category.	P Document published on or after the declared priority date but before the filing date of this invention.
& Member of the same patent family	E Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC*:

G4A

Worldwide search of patent documents classified in the following areas of the IPC*:

G06F, I105K

The following online and other databases have been used in the preparation of this search report:

WPI, EPODOC, PAJ, INSPEC, IBM TDB, IEEE Xplore, Internet